

Handling Qualities Evaluations of Low Complexity Model Reference Adaptive Controllers for Reduced Pitch and Roll Damping Scenarios

Curt Hanson¹, Jacob Schaefer², and John J. Burken³
NASA Dryden Flight Research Center, Edwards, CA 93523

Marcus Johnson⁴
TYBRIN Corporation, Edwards, CA 93523

Nhan Nguyen⁵
NASA Ames Research Center, Mountain View, CA 94035

National Aeronautics and Space Administration (NASA) researchers have conducted a series of flight experiments designed to study the effects of varying levels of adaptive controller complexity on the performance and handling qualities of an aircraft under various simulated failure or damage conditions. A baseline, nonlinear dynamic inversion controller was augmented with three variations of a model reference adaptive control design. The simplest design consisted of a single adaptive parameter in each of the pitch and roll axes computed using a basic gradient-based update law. A second design was built upon the first by increasing the complexity of the update law. The third and most complex design added an additional adaptive parameter to each axis. Flight tests were conducted using NASA's Full-scale Advanced Systems Testbed, a highly modified F-18 aircraft that contains a research flight control system capable of housing advanced flight controls experiments. Each controller was evaluated against a suite of simulated failures and damage ranging from destabilization of the pitch and roll axes to significant coupling between the axes. Two pilots evaluated the three adaptive controllers as well as the non-adaptive baseline controller in a variety of dynamic maneuvers and precision flying tasks designed to uncover potential deficiencies in the handling qualities of the aircraft, and adverse interactions between the pilot and the adaptive controllers. The work was completed as part of the Integrated Resilient Aircraft Control Project under NASA's Aviation Safety Program.

Nomenclature

Acronyms

CH	= Cooper-Harper
FAST	= Full-scale Advanced Systems Testbed
HQ	= handling qualities
HQR	= handling qualities rating
HUD	= head up display
IFCS	= intelligent flight control system
IRAC	= Integrated Resilient Aircraft Control
KIAS	= knots indicated airspeed
MRAC	= model reference adaptive control

¹ Aerospace Engineer, Flight Controls Branch, P.O. Box 273/Mail Stop 4840D, AIAA Member.

² Aerospace Engineer, Flight Controls Branch, P.O. Box 273/Mail Stop 4840D, AIAA Member.

³ Aerospace Engineer, Flight Controls Branch, P.O. Box 273/Mail Stop 4840D, AIAA Senior Member.

⁴ Aerospace Engineer, Flight Controls Branch, P.O. Box 273/Mail Stop 4840D, AIAA Member.

⁵ Research Scientist, Intelligent Systems Division, Mail Stop 269-1, AIAA Associate Fellow.

NDI	= nonlinear dynamic inversion
OCM	= optimal control modification
onMRAC	= MRAC with normalization and OCM
onMRAC+	= MRAC with normalization, OCM, and adaptation for time-varying disturbances
PI	= proportional-plus-integral
PIO	= pilot in the loop oscillation
RFI	= request for information
SDC	= self designing controller
sMRAC	= simple MRAC

Mathematical Symbols

A	= matrix of state derivative coefficients
B	= matrix of input coefficients
C	= NDI error compensator transfer function
C_{l_p}	= roll damping coefficient
C_{m_q}	= pitch damping coefficient
f_A	= aerodynamic and inertial moments
g	= load factor
H	= vector of angular momentum
I	= inertia matrix
k	= reference model input gain
L_α	= apparent lift curve slope
M	= applied moments
N	= normalization weighting matrix
P	= solution to the Lyapunov equation
p	= roll rate, degrees per second
Q	= tunable constant matrix of the Lyapunov equation
q	= pitch rate, degrees per second
R	= reference model
r	= reference input, yaw rate
s	= Laplace operator
t	= time, seconds
u	= scalar control input
W	= pseudo-inverse weighting matrix
x	= state vector
\tilde{x}	= MRAC reference model tracking error
y	= vector of system outputs
α	= angle of attack
Γ	= adaptive parameter learning rate
δ	= control surface deflections, pilot control inputs
ζ	= damping ratio
Θ	= vector of matched uncertainties
$\hat{\Theta}$	= vector of matched uncertainty estimates
θ	= matched uncertainty
$\hat{\theta}$	= matched uncertainty estimate
ν	= OCM tunable parameter
σ	= scalar, time-varying disturbance uncertainty
$\hat{\sigma}$	= scalar, time-varying disturbance uncertainty estimate
Ω	= vector of rotational rates
ω	= reference model frequency
\mathbb{R}	= set of real numbers

Subscripts

0	= trim
<i>a</i>	= adaptive control
<i>c</i>	= NDI error compensator
<i>cmd</i>	= command
<i>err</i>	= NDI reference model tracking error
<i>l</i>	= rolling moment
<i>lat</i>	= lateral
<i>lon</i>	= longitudinal
<i>m</i>	= reference model, pitching moment
<i>max</i>	= maximum achievable yaw rate
<i>p</i>	= roll axis
<i>q</i>	= pitch axis
<i>q</i> ₁	= pitch axis adaptive parameter associated with the integral of pitch rate
<i>q</i> ₂	= pitch axis adaptive parameter associated with pitch rate
<i>ref</i>	= NDI reference model
α	= angle of attack
δ	= control surfaces

Superscripts

<i>m</i>	= number of control effectors
<i>n</i>	= order of the reference model

I. Introduction

Flight research of adaptive controls has been sporadic over the past 50 years leaving relatively little flight data on their application to full-scale piloted aircraft to guide present-day designers. Examples that do exist include the first flight tests of a model reference adaptive system on the F-94A aircraft (Lockheed Martin, Bethesda, Maryland),¹ an experimental adaptive flight control system evaluated on the F-94C aircraft (Lockheed Martin, Bethesda, Maryland),² the implementation of an adaptive flight control system on the X-15 aircraft (North American Aviation, Inc., Downy, California),^{3,4} testing of an indirect-adaptive self designing controller (SDC) on the F-16 VISTA (General Dynamics, now Lockheed Martin, Bethesda, Maryland)⁵ and the intelligent flight control system (IFCS) research on the highly-modified NASA F-15 aircraft (McDonnell Douglas, now The Boeing Company).^{6,7} By the mid-2000s, experiments such as SDC and IFCS had demonstrated that adaptive flight controls can be an effective technology for improving aircraft safety in the event of failures or damage. However, the nonlinear, time-varying nature of adaptive systems continues to challenge traditional methods of verification and validation of safety-critical flight control systems.

In April, 2009 the National Aeronautics and Space Administration's Integrated Resilient Aircraft Control (IRAC) project disseminated a request for information (RFI) to the adaptive controls community seeking ideas for potential flight experiments.⁸ A workshop was held in Chicago in August of 2009 with representatives from industry, academia, and other government agencies to discuss the wide variety of RFI responses received by NASA. Three focus areas were identified through this process.

1. Simple, yet effective, adaptive control algorithms should be investigated to help address the issue of verification and validation of adaptive flight controls.
2. The appropriate level of pilot awareness and interaction with adaptive control systems should be studied, including the potential for adverse interactions such as pilot-in-the-loop oscillations (PIO).
3. Techniques should be matured for incorporating feedback information, both static and dynamic, from an aircraft's structure into the flight control system.

The first, and to some extent, the second of these focus areas prompted the application of a low complexity, textbook-like direct model reference adaptive control (MRAC) scheme to the NASA Full-scale Advanced Systems Testbed (FAST). FAST is a highly modified F-18 aircraft (Lockheed Martin, Bethesda, Maryland) that contains a research flight control system capable of housing advanced flight controls experiments. Flight experimentation of adaptive systems using a piloted, full-scale testbed has been shown to uncover implementation issues that may not be found through experiments with simulations or sub-scale, remotely piloted aircraft. This report discusses the results of flight research with several variations of the low complexity MRAC design. The controllers were

evaluated for a healthy aircraft and two simulated failure scenarios, characterized by severe loss of damping in either the pitch or the roll axis.

A philosophy of “simpler is better” motivated the development of the low complexity MRAC formulation under the assumption that simplification leads to lower implementation and verification costs, and ultimately to a greater likelihood of acceptance by aircraft manufacturers, operators and certification authorities. Reduced complexity also tends to produce a safer implementation by reducing the potential for implementation errors within the adaptive controller software.

Evaluation of the relative complexity of various controls techniques, including adaptive control methods, requires ways of measuring complexity. Several aspects of an adaptive controller were identified as having impacts to the complexity of the design, implementation, and testing process including flight test. Simplification was achieved by minimizing these sources of complexity.

1. Complex mathematical elements are mathematical constructs, such as integrators and filters, which require initialization or internal state limiting, whose implementation requires special attention or care during design and testing. Other examples include division and inverse functions, which must be protected from divide-by-zero scenarios.
2. Free design parameters such as gains, learning rates, and filter coefficients are elements of the controller that must be tuned to achieve the desired combination of controller performance and robustness. Increasing the number of free design parameters tends to complicate the design and verification effort of the controller due to the potential for unforeseen interactions between them and the increased difficulty of intuitively identifying cause and effect.
3. Input / output plane dimensionality of the adaptive controller impacts control system analysis and testing. Robustness analysis is typically performed at each feedback loop and each output. Noise, latency, and sampling rate characteristics must be evaluated and analyzed for each input, and actuator dynamics for each output. During development and testing, the failure mode of each input and output must be characterized, and their effects evaluated.

Three variations of a model reference adaptive control design were implemented to augment a baseline, nonlinear dynamic inversion (NDI) controller. The simplest design consisted of a single adaptive parameter in each of the pitch and roll axes computed using a basic gradient-based update law. A second design increased the complexity of the update law, while the third added another adaptive parameter in each axis.

II. General Description of Model Reference Adaptive Control

Model reference adaptive control derives its name from the use of reference model dynamics to define a desired trajectory for the system outputs to follow.⁹ The reference model for the aircraft axis of interest is specified as a stable, linear time-invariant system.

$$\dot{x}_m(t) = A_m x_m(t) + B_m r(t) \quad (1)$$

The scalar reference input $r(t)$ in Eq. (1) is a function of the time-varying pilot control command. Constant matrices $A_m \in \mathbb{R}^{n \times n}$ and $B_m \in \mathbb{R}^{n \times 1}$ are selected to give good aircraft handling qualities. The aircraft's true dynamics are written as Eq. (2),

$$\dot{x}(t) = Ax(t) + B(u(t) - \sigma(t)) \quad (2)$$

where A and B are unknown and $\sigma(t)$ is a scalar, time-varying uncertain disturbance. The unknown aircraft dynamics and known reference model dynamics are related according to the usual MRAC matching conditions⁹ in Eqs. (3) and (4),

$$A = A_m - B\Theta^T \text{ and} \quad (3)$$

$$B = B_m \quad (4)$$

Several restrictions are apparent from the matching conditions. The order of the reference model n must match the order of the aircraft dynamics, which also defines the number of matched uncertainties in the parameter vector Θ . These uncertainties, commonly referred to as A-matrix uncertainties, parameterize changes to the pitch and roll

damping characteristics of the aircraft. It is assumed that these uncertainties are either constant or that they vary slowly in comparison to the aircraft's dynamics.

The scalar control input $u(t)$ is computed as a combination of the scalar reference input and augmentation from the adaptive controller as shown in Eq. (5).

$$u(t) = r(t) + \dot{x}_a(t) = r(t) + \hat{\Theta}^T(t)x(t) + \hat{\sigma}(t) \quad (5)$$

$\hat{\Theta}(t)$ and $\hat{\sigma}(t)$ are adaptive parameters to be estimated. The control input from Eq. (5) is substituted into the aircraft dynamics of Eq. (2) along with the matching conditions of Eqs. (3) and (4), to find Eq. (6),

$$\dot{x}(t) = A_m x(t) + B_m r(t) + B_m (\hat{\Theta}^T(t) - \Theta^T) x(t) + B_m (\hat{\sigma}(t) - \sigma(t)). \quad (6)$$

As adaptation takes place, $\hat{\Theta}(t) \rightarrow \Theta$, $\hat{\sigma}(t) \rightarrow \sigma$, and the aircraft dynamics in Eq. (6) approach those of the reference model in Eq. (1). The aircraft dynamics become Eq. (7),

$$\dot{x}(t) = A_m x(t) + B_m r(t) \quad (7)$$

III. Implementation of MRAC with Nonlinear Dynamic Inversion

The adaptive controller is implemented as a pilot-selectable augmentation mode to a non-adaptive baseline controller. This approach is representative of a pilot-initiated emergency system. Adaptive augmentation is applied only to the pitch and roll axes of the aircraft, as the characteristics of these axes tend to dominate the pilot's perception of handling qualities.

The non-adaptive baseline is a full, NDI inner-loop controller. The NDI inner loop includes a proportional-plus-integral (PI) output error feedback compensator for improved robustness. NDI was chosen for its analyzability, open-source architecture, and convenient methods of introducing simulated aircraft failures and damage through its explicit model-following architecture and on-board aerodynamic lookup tables.

The rotational equations of motion for the FAST aircraft can be written using the relationship between the rate of change of momentum and applied moments under the assumption of constant inertias¹⁰ as shown in Eq. (8):

$$\frac{d}{dt} H(t) = I\dot{\Omega}(t) + \Omega(t) \times I\Omega(t) = \sum M(t) \quad (8)$$

Equation (8) is solved for the rotational acceleration vector $I\dot{\Omega}(t)$ as a function of the inertial coupling terms $\Omega(t) \times I\Omega(t)$ and the applied moments $\sum M(t)$. In Eq. (9), the applied moments are expanded into the applied aerodynamic and inertial moments f_A , the applied moments due to control surface deflections $\delta(t) - \delta_0$, and the applied moments due to disturbances $\sigma(t)$.

$$I\dot{\Omega}(t) = f_A(y(t)) + B_\delta(\delta(t) - \delta_0) - (\Omega(t) \times I\Omega(t)) - I\sigma(t) \quad (9)$$

Here, $y(t)$ is the output vector consisting of rotational body rates, and the aerodynamic angles of attack and sideslip. The rotational acceleration $\dot{\Omega}(t)$ on the left-hand side is replaced by the desired rotational acceleration command, $\dot{x}_{cmd}(t)$, and it is assumed that $\sigma(t) = 0$. Equation (9) is then solved for the control surface deflection vector $\delta(t)$ necessary to achieve the commanded dynamics as shown in Eq. (10).

$$\delta(t) = B_\delta^{-1} (I\dot{x}_{cmd}(t) + \Omega(t) \times I\Omega(t) - \hat{f}_A(y(t))) + \delta_0 \quad (10)$$

$\hat{f}_A(y(t))$ is an estimate computed from on-board aerodynamic lookup tables. Since $B_\delta \in \mathbb{R}^{3 \times m}$, $m > 3$, the weighted pseudo-inverse shown in Eq. (11) must be used to compute its inverse.

$$B_\delta^{-1} = W^{-1} B_\delta^T (B_\delta W^{-1} B_\delta^T)^{-1} \quad (11)$$

Additional details on the NDI implementation, including the calculation of $\hat{f}_A(y(t))$, B_δ and the trim control state δ_0 can be found in Miller.¹¹ Figure 1 shows the top-level block diagram of the integrated NDI and MRAC controllers. The NDI controls all three axes, although adaptive control is only implemented in the pitch and roll axes. Equation (12) gives the expanded acceleration command vector for each axis.

$$\dot{x}_{cmd}(t) = \dot{x}_{ref}(t) + \dot{x}_c(t) + \dot{x}_a(t) = \begin{bmatrix} \dot{p}_{ref}(t) + \dot{p}_c(t) + \dot{p}_a(t) \\ \dot{q}_{ref}(t) + \dot{q}_c(t) + \dot{q}_a(t) \\ \dot{r}_{ref}(t) + \dot{r}_c(t) \end{bmatrix} \quad (12)$$

Note that, in Eq. (12), the $\dot{x}_c(t)$ terms are the outputs of the error compensator C of Fig. 1, and the $\dot{x}_a(t)$ terms are the adaptive control contributions. The implementation details differ slightly between the pitch and the roll axes. Each axis will be discussed individually.

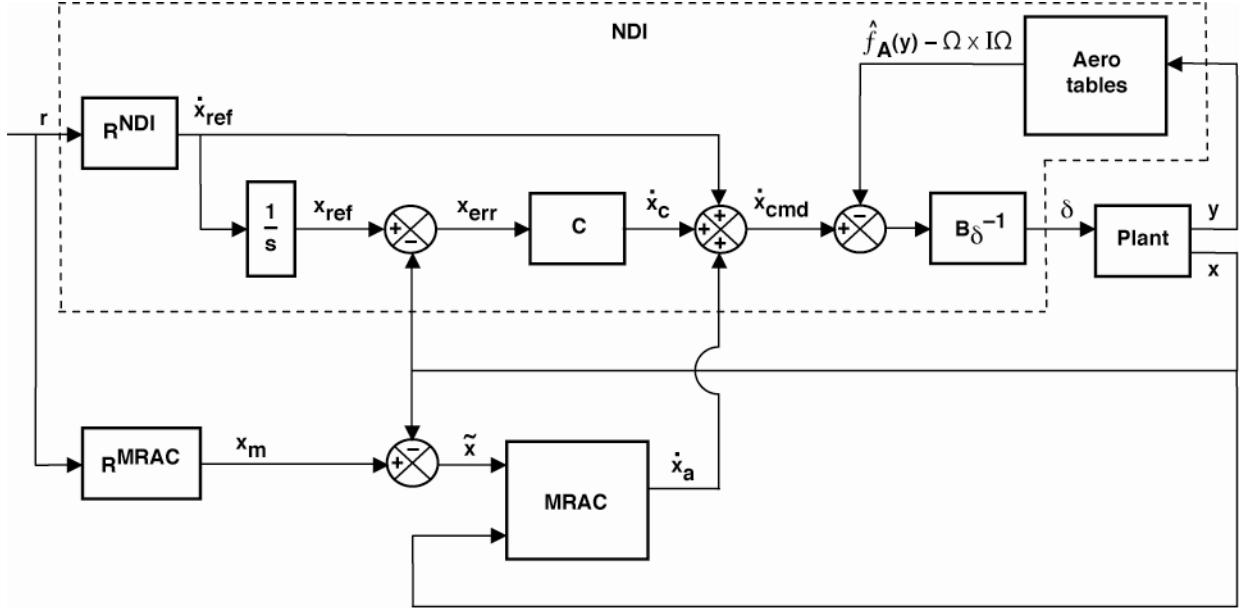


Figure 1. MRAC with NDI block diagram.

A. Pitch Axis Implementation

Reference models for the NDI and MRAC are implemented separately to allow the NDI to simulate the dynamics of a damaged aircraft. Both use the second-order transfer function shown in Eq. (13) to model the aircraft short period dynamics in the pitch axis.¹²

$$\frac{\dot{q}_m(s)}{r_q(s)} = \frac{s^2}{s^2 + 2\zeta_q\omega_q s + \omega_q^2} \quad (13)$$

The pitch reference input signal is a function of the pilot longitudinal stick command δ_{lon} described by Eq. (14).

$$\frac{r_q(s)}{\delta_{lon}(s)} = \frac{k_q\omega_q^2(s + L_\alpha)}{s} \quad (14)$$

An error compensator is necessary to improve tracking of the reference model in the presence of model uncertainty and disturbances by the NDI controller without adaptive augmentation. The NDI pitch axis contains a proportional plus integral (PI) error compensator. The gains of the compensator in Eq. (15) are tuned to match the denominator of the NDI reference model in Eq. (13).

$$C_q(s) = \frac{\dot{q}_c(s)}{q_{err}(s)} = \frac{2\zeta_q \omega_q s + \omega_q^2}{s} \quad (15)$$

The pitch rate error signal is calculated as the difference between the output of the NDI reference model $q_{ref}(t)$ and the pitch rate feedback as shown in Eq. (16).

$$q_{err}(t) = q_{ref}(t) - q(t) \quad (16)$$

The adaptive control command is computed as shown in Eq. (17).

$$\dot{q}_a(t) = \hat{\theta}_{q1}(t) \int q(t) + \hat{\theta}_{q2} q(t) + \hat{\sigma}_q(t) \quad (17)$$

B. Roll Axis Implementation

The first-order transfer function shown in Eq. (18) is used as the reference model of the stability axis roll dynamics.¹¹

$$\frac{\dot{p}_m(s)}{r_p(s)} = \frac{s}{s + \omega_p} \quad (18)$$

The roll reference input signal is a function of the pilot lateral stick command δ_{lat} described by Eq. (19).

$$r_p(t) = \left(\frac{r_{max}}{\alpha} \right) k_p \omega_p \delta_{lat}(t) \quad (19)$$

The NDI roll axis contains a proportional error compensator whose gains are tuned to the denominator of the NDI reference model. Note that for handling qualities, the roll rate feedback is converted to the stability axis prior to computing its error with the NDI reference model command¹¹ as shown in Eq. (20).

$$\dot{p}_c(t) = \omega_p (p_{ref}(t) - p(t)) \quad (20)$$

Since the roll rate transfer function is first-order, it follows that the roll axis adaptive augmentation contains a single matched uncertainty term as shown in Eq. (21).

$$\dot{p}_a(t) = \hat{\theta}_p(t) p(t) + \hat{\sigma}_p(t) \quad (21)$$

IV. MRAC Design

Three update laws were designed with increasing levels of complexity. The free design parameters of the update laws were tuned in the piloted simulation for acceptable response when no failures were present as well as across a wide variety of simulated failures and damage scenarios. Aggressive pilot inputs were used during tuning to stress the system, and parameters were adjusted to achieve the fastest adaptation possible without introducing undesirable side effects such as actuator rate limiting.

A. sMRAC Update Laws

Simple gradient-based update laws for the pitch and roll adaptive parameters is given in Eqs. (22) and (23). These update laws are referred to as simple MRAC, or sMRAC for short.

$$\begin{bmatrix} \dot{\hat{\theta}}_{q1} \\ \dot{\hat{\theta}}_{q2} \end{bmatrix} = \begin{bmatrix} \Gamma_{q1} & 0 \\ 0 & \Gamma_{q2} \end{bmatrix} \tilde{x}_q^T P_q B_{m_q} x_q \quad (22)$$

$$\dot{\hat{\theta}}_p = \Gamma_p \tilde{x}_p^T P_p B_{m_p} x_p \quad (23)$$

Γ_{q1} , Γ_{q2} , and Γ_p are tunable, positive scalar learning rates. Tracking error is computed by subtracting the aircraft response from the desired dynamics of the reference model as shown in Eqs. (24) and (25).

$$\tilde{x}_q(t) = x_{q_m} - x_q = \begin{bmatrix} \int q_m(t) - \int q(t) \\ q_m(t) - q(t) \end{bmatrix} \quad (24)$$

$$\tilde{x}_p(t) = x_{p_m} - x_p = p_m(t) - p(t) \quad (25)$$

The pitch axis gain matrix $P_q = P_q^T$ is a positive definite solution to the Lyapunov equation in Eq. (26), where $Q_q = Q_q^T$ is also positive definite, selected by the designer. The roll axis gain matrix P_p is computed in a similar manner.

$$P_q A_{m_q} + A_{m_q}^T P_q = -Q_q \quad (26)$$

In the pitch axis, the influence of $\int q_m(t)$ and $\int q(t)$ is minimized through tuning of the Γ_{q_1} and Q_q matrices. This leaves essentially a single adaptive parameter in both the pitch and roll axes of the sMRAC controller.

Initial simulation evaluations of the update laws in Eqs. (22) and (23) showed that the sMRAC controller adapted very well to cancel matched uncertainties, but performed poorly in cases where the applied uncertainty was not well-parameterized by the matched uncertainty estimates, or when rapid commands were generated by the pilot. In these cases, the adaptive controller did not completely eliminate the tracking error, and consequently, over-adapted to the point where control surface rate limiting was encountered and the closed-loop system became unstable.

B. onMRAC Update Laws

The additional complexity of an update law modification term was added, along with a normalization term, to counter the undesirable tendency observed in piloted simulation of the sMRAC update law to over-adapt. An optimal control modification (OCM) term was selected.¹³ The resulting update law, referred to as optimal normalized MRAC, or onMRAC for short, is described by Eqs. (27) and (28).

$$\begin{bmatrix} \dot{\hat{\theta}}_{q_1} \\ \dot{\hat{\theta}}_{q_2} \end{bmatrix} = \frac{\begin{bmatrix} \Gamma_{q_1} & 0 \\ 0 & \Gamma_{q_2} \end{bmatrix}}{1 + x_q^T N_{\theta_q} x_q} \left(\tilde{x}_q^T P_q B_{m_q} x_q - x_q \nu_{\theta_q} x_q^T \begin{bmatrix} \hat{\theta}_{q_1} \\ \hat{\theta}_{q_2} \end{bmatrix} B_{m_q}^T P_q A_{m_q}^{-1} B_{m_q} \right) \quad (27)$$

$$\dot{\hat{\theta}}_p = \frac{\Gamma_p}{1 + x_p^T N_{\theta_p} x_p} \left(\tilde{x}_p^T P_p B_{m_p} x_p - x_p \nu_{\theta_p} x_p^T \hat{\theta}_p B_{m_p}^T P_p A_{m_p}^{-1} B_{m_p} \right) \quad (28)$$

The designer-selected tuning gains ν_{θ_q} and ν_{θ_p} control the relative influence that the OCM terms have to dampen oscillations in the adaptive parameters and to suppress over-adaptation. The weighting matrices N_{θ_q} and N_{θ_p} control the extent to which the rate of adaptation is reduced during large aircraft dynamics.

C. onMRAC+ Update Laws

A further level of complexity was explored through the inclusion of an additional adaptive parameter in each axis, $\hat{\sigma}_q$ and $\hat{\sigma}_p$, corresponding to uncertain disturbances in the pitch and roll axis, respectively. These adaptive parameter differ from $\hat{\theta}_{q_1}$ and $\hat{\theta}_p$ in that they do not rely upon the basis functions x_q and x_p . This configuration is referred to as the optimal normalized MRAC plus, or onMRAC+ for short, and described by Eqs. (29) and (30).

$$\dot{\hat{\sigma}}_q = \frac{\Gamma_{\sigma_q}}{1 + x_q^T R_{\sigma_q} x_q} \left(\tilde{x}_q^T P_q B_{m_q} - \nu_{\sigma_q} \hat{\sigma}_q B_{m_q}^T P_q A_{m_q}^{-1} B_{m_q} \right) \quad (29)$$

$$\dot{\hat{\sigma}}_p = \frac{\Gamma_{\sigma_p}}{1 + x_p^T R_{\sigma_p} x_p} \left(\tilde{x}_p^T P_p B_{m_p} - \nu_{\sigma_p} \hat{\sigma}_p B_{m_p}^T P_p A_{m_p}^{-1} B_{m_p} \right) \quad (30)$$

In the pitch axis, the OCM tuning parameter v_{σ_q} was zeroed out to aid in the investigation of the usefulness of the OCM term. Inclusion of normalization without the OCM term in the pitch axis contrasts with the inclusion of both terms in the roll axis and provides insight into whether normalization alone is sufficient to ensure desirable adaptive parameter characteristics.

V. Flight Research Methodology

The controllers were evaluated by their ability to adapt to simulated in-flight failures/damage rather than to changes in flight condition. Therefore, the controllers were tuned for a single flight condition of 240 KIAS and 25,000 ft pressure altitude. This airspeed-altitude combination falls within the envelope cleared for the FAST aircraft in which any unexpected control surface deflection combination will not exceed structural load limits and ensures that there is adequate altitude for recovery in the case of a departure from controlled flight.

Initial evaluations of the baseline NDI controller and the three adaptive controllers were completed for a healthy aircraft as well as for an 80% reduction in the pitch damping coefficient C_{m_q} and a 117% reduction (resulting in a change in sign) in the roll damping coefficient C_{l_p} . The reduced pitch and roll damping failures were implemented by altering the values of C_{m_q} and C_{l_p} in the NDI aerodynamic lookup tables. This approach effectively made use of the NDI controller as a mechanism to model a failed aircraft. These coefficients map very closely to the matched uncertainties θ_{q2} and θ_p . Additional failure scenarios were tested and are described in Ref. 14.

A progression of standard flight test clearance maneuvers, shown below, was designed to investigate the basic flying qualities of the system and clear the controllers for handling qualities evaluations. Pilot inputs, aircraft states, control surface positions, and internal controller parameters were recorded on-board and telemetered for real-time and post-flight analysis.

1. 3-axis doublets
2. pitch and bank angle captures
3. steady-heading sideslips to maximum rudder pedal deflection
4. 360-degree rolls at half-stick deflection
5. 2.5 g windup turns
6. 2 g loaded roll reversals through 120 degrees at half-stick deflection

Following completion of the clearance maneuvers, piloted handling quality tasks were completed for each controller with four of the five simulated failures. The reduced pitch damping failure was evaluated in a 2 g air-to-air tracking task. An in-trail formation tracking task was used for the roll damping failure. Both tasks were also performed without any simulated failures to provide a baseline for comparison. Envelope restrictions on the research flight control system currently prevent the execution of tasks such as approach and landing with the experimental control laws active. However, there is some evidence that an in-trail tracking task is a good approximation of an offset landing task.¹⁵

The 2 g air-to-air tracking task was accomplished using the head-up-display (HUD) air-to-ground reticle set to a deflection of 140 mil. Aircraft spacing was approximately 1,000 ft. The task began with an initial 2 g turn by both aircraft, followed by a call from the test aircraft for the target to reverse the direction of its turn. Upon reversal, the pilot of the test aircraft completed a gross acquisition task, followed by a fine tracking task. The desired criterion for the gross acquisition task was placement of the target aircraft within the reticle with no overshoots. Gross acquisition with a maximum of one overshoot was considered adequate. Desired fine tracking performance required placement of the reticle's center pipper on the target aircraft center of mass at least 50% of the time. Placement of the pipper anywhere on the target for a minimum of 50% of the task duration was considered adequate.

The in-trail formation tracking task was accomplished in straight and level flight. The test aircraft was initially positioned directly behind and slightly below the target aircraft using a mark on the canopy to achieve a consistent starting point. Nose-to-tail spacing was approximately two ship lengths. The gross acquisition task required the pilot to execute a curvilinear trajectory from the centerline starting position to a point where the top inboard corner of the HUD bracket aligned with the wingtip of the target while keeping the nose-to-tail spacing constant. The fine tracking task then required the pilot to maintain this position. Desired performance in the execution of the gross acquisition task required no more than one overshoot, while adequate performance allowed at most two overshoots. Desired fine tracking required the pilot to maintain the HUD bracket on the wingtip for no less than 80% of the duration of the task. A criterion of at least 50% was specified for adequate fine tracking performance.

VI. Flight Results

Handling qualities evaluations of the controllers with reduced pitch and roll damping failure scenarios as well as with a healthy aircraft are discussed in this section. An analysis of the failure scenarios exhibiting cross-coupling characteristics is presented in the companion paper by Schaefer.¹⁴

A. Healthy Aircraft Scenario

Three NASA test pilots, designated by the letters A through C, participated in the experiment. Cooper-Harper (CH)¹⁶ handling quality ratings are shown in Fig. 2 for the four controllers evaluated in a 2 g tracking task with no simulated failures. Only one pilot, Pilot A, evaluated all four controllers in this scenario while Pilot B and Pilot C evaluated only the NDI controller. The ratings indicate that the baseline, non-adaptive controller has reasonably good handling qualities, rated as satisfactory without improvement by all of the pilots with the exception of Pilot B's assessment of fine tracking, which he commented on as being "a little pitch sensitive."

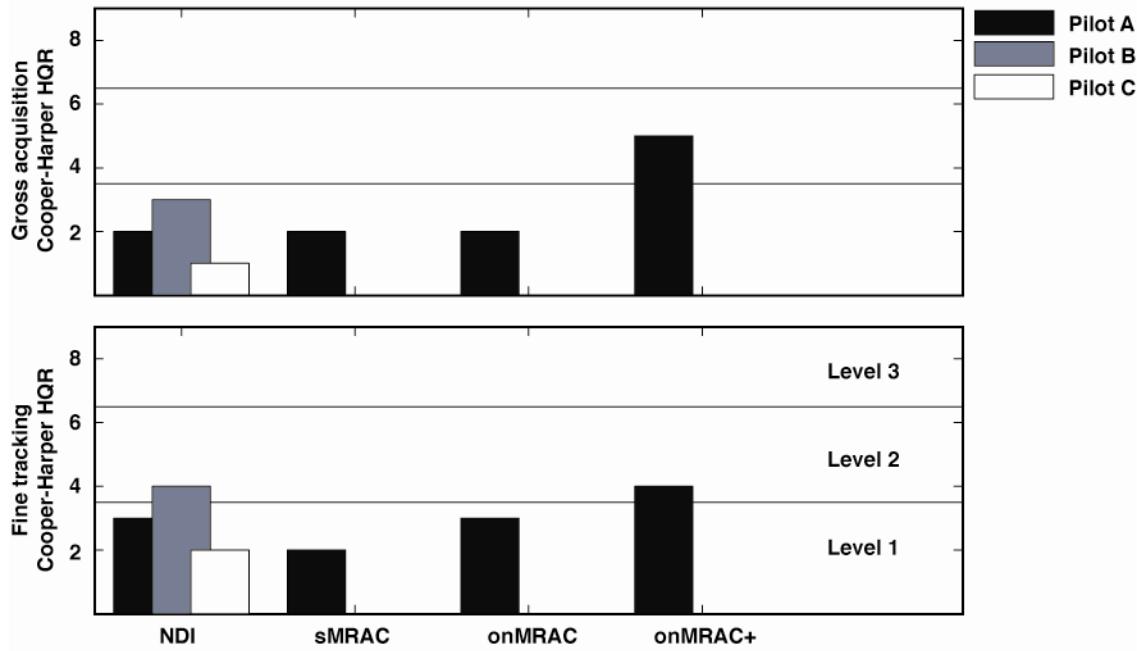


Figure 2. Handling qualities ratings for a 2 g tracking task with no failure.

A design goal for the adaptive controllers was to ensure that they did not significantly degrade the handling qualities of the aircraft in the un-failed state; which was referred to as the "do no harm" requirement. Pilot A's assessment of the sMRAC and onMRAC controllers indicate that they maintained desired Level 1 flying qualities. The onMRAC+ controller was rated significantly worse, particularly in gross acquisition with a CH rating of 5 indicating moderately objectionable deficiencies. Pilot A commented on the pitch axis having "a lot more bobble in gross acquisition than I would have expected; seemed like there was always like 2 overshoots in the gross acquisition."

A common metric for evaluating the performance of model-reference adaptive controllers is the integral of the absolute value of the tracking error between the MRAC reference model and aircraft dynamics.¹⁷ Figure 3 shows the pitch rate tracking error metric during both the gross acquisition and fine tracking tasks performed by Pilot A for each controller. The pitch tracking error for the adaptive controllers during gross acquisition was only slightly worse than that of the NDI while the fine tracking error was nearly identical for all four controllers. It should also be noted that small differences in tracking error are not necessarily indicative of variations in controller performance, but may instead reflect differences in pilot task execution. Therefore, an analysis of tracking performance alone is inconclusive with respect to providing a satisfactory explanation for the degradation in pilot ratings for the onMRAC+ controller.

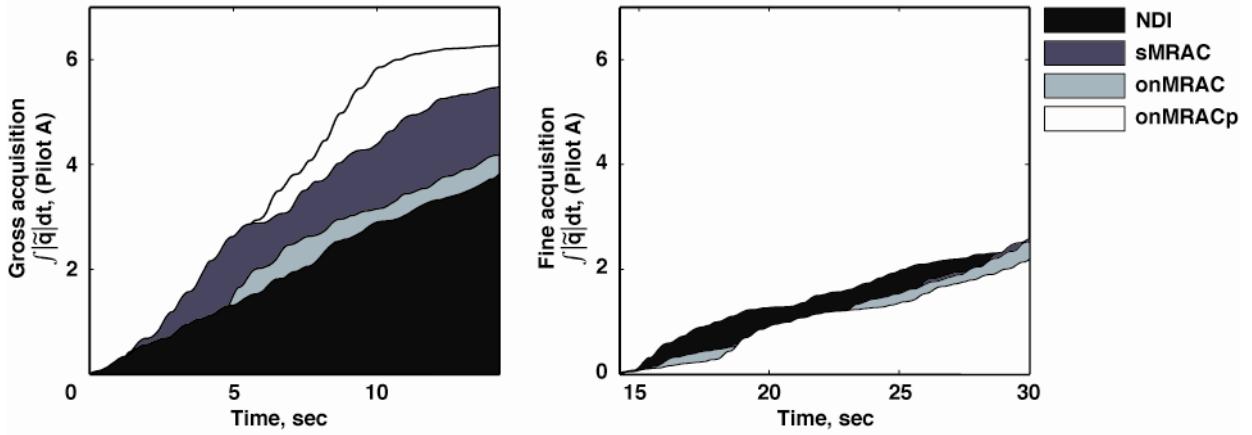


Figure 3. Pitch rate tracking error for a 2 g tracking task with no failure (Pilot A).

A clearer understanding of the behavior of the adaptive controllers can be found by examining the adaptive parameters in Fig. 4 for a sample 2 g tracking maneuver by Pilot A. Because there is no simulated failure, the only uncertainty is due to a small amount of error in the dynamic inversion,¹⁸ and consequently the ideal adaptive parameter values are close to zero. Throughout the maneuver, values of the matched uncertainty parameters for the sMRAC and onMRAC controllers remain small; recall that neither of these controllers has a $\hat{\sigma}_q$ term. The onMRAC+ controller shows an oscillatory tendency in the $\hat{\sigma}_q$ parameter, particularly during the gross acquisition task. It is shown in Ref. 19 that in general, the $\hat{\sigma}_q$ parameter is adapting to input uncertainty in the signal r_q as well as to counter external disturbances. During the 2 g tracking task, there was no indication of an encounter with the jet wash of the target aircraft or other atmospheric disturbance. Variations in $\hat{\sigma}_q$ therefore are directly altering the pilot's pitch command path, introducing a lack of predictability (from the pilot's perspective) in the pitch responsiveness of the aircraft. The result is an adverse interaction between the pilot and the adaptive controller, which evolves into a minor pilot-in-the-loop oscillation (PIO) and triggers overly aggressive inputs. Pilot A gave a PIO rating of 3 for this controller in the gross acquisition task indicating that undesirable motions tended to occur following abrupt maneuvers. A comment by the pilot indicated that he "had to back out of the loop a little bit when doing the gross acquisition, let everything kind of die out."

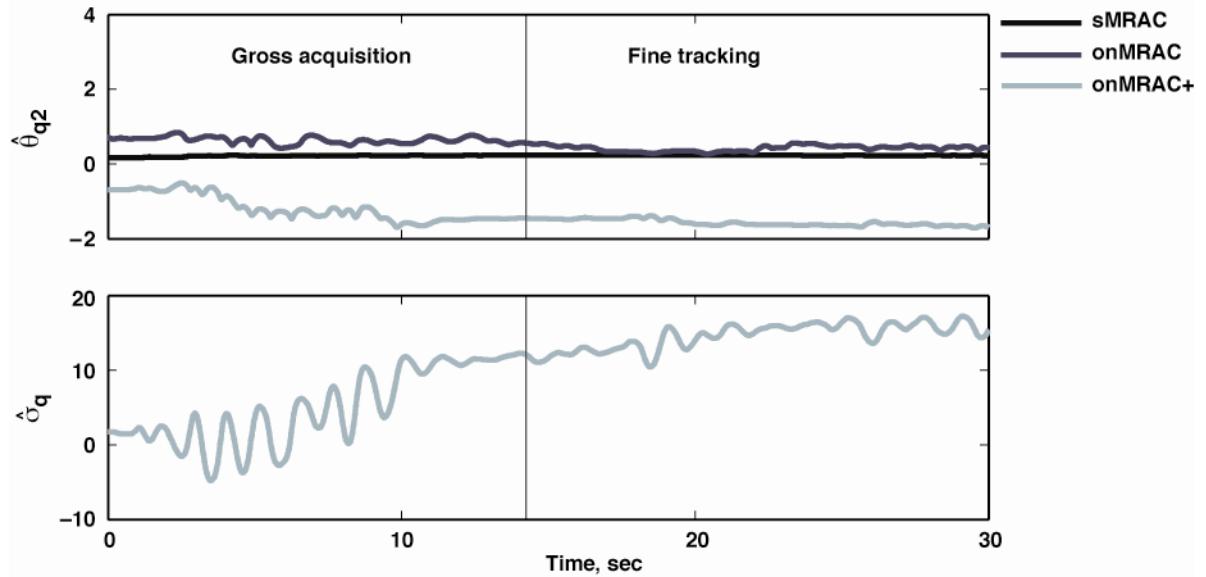


Figure 4. Adaptive parameter values for a 2 g tracking task with no failure (Pilot A).

Also from Fig. 4, $\hat{\theta}_{q2}$ is more negative for onMRAC+ than for the other controllers, caused by tracking errors generated during the aggressive inputs during the gross acquisition task. A negative value for $\hat{\theta}_{q2}$ has the effect of increased damping in the pitch response of the aircraft, making it less PIO-prone.

B. Reduced Pitch Damping Failure

Two pilots evaluated the reduced pitch damping failure during a 2 g air-to-air tracking task. Handling qualities ratings given by Pilot A and Pilot B for each of the four controllers during the 2 g air-to-air tracking task with a reduced pitch damping failure are tabulated in Fig. 5. Due to the high workload required to maintain pitch control in gross acquisition, Pilot B was unable to maintain sufficient airspeed to complete the fine tracking portion of the task with the NDI controller, and consequently no rating was given for that case. An overall assessment of the ratings indicate that all three of the adaptive controllers improved the handling qualities of the aircraft during the gross acquisition portion of the task, but that the varying levels of controller complexity had minimal effect on the ratings. Little to no improvement was observed in the ratings for the fine tracking portion of the task even when compared against the NDI controller without adaptation.

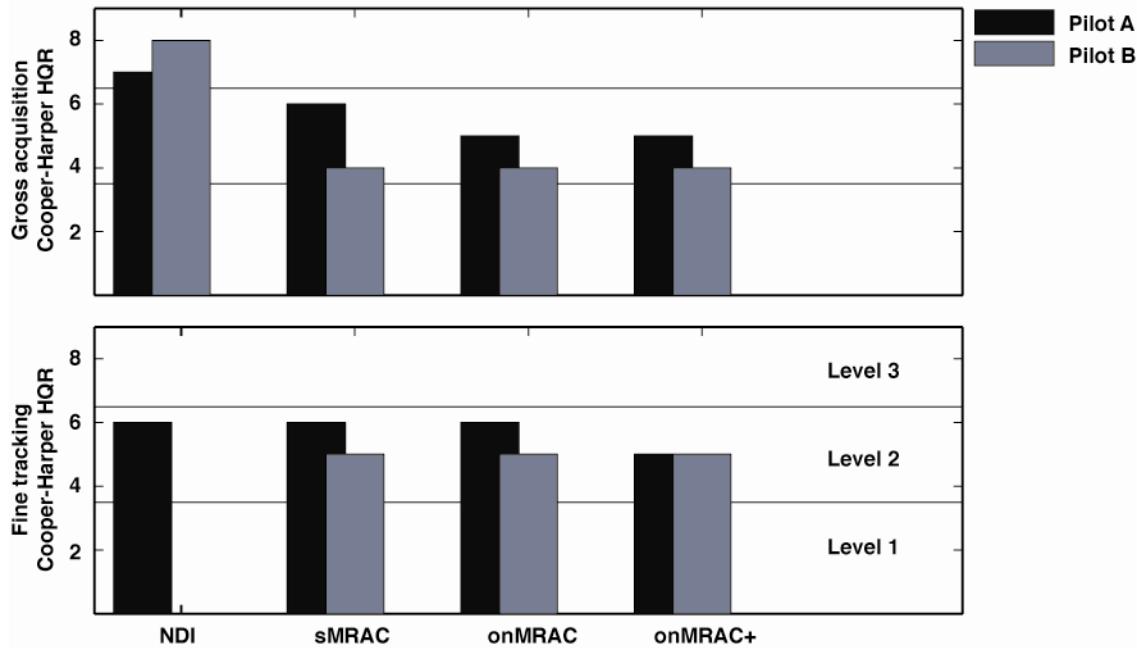
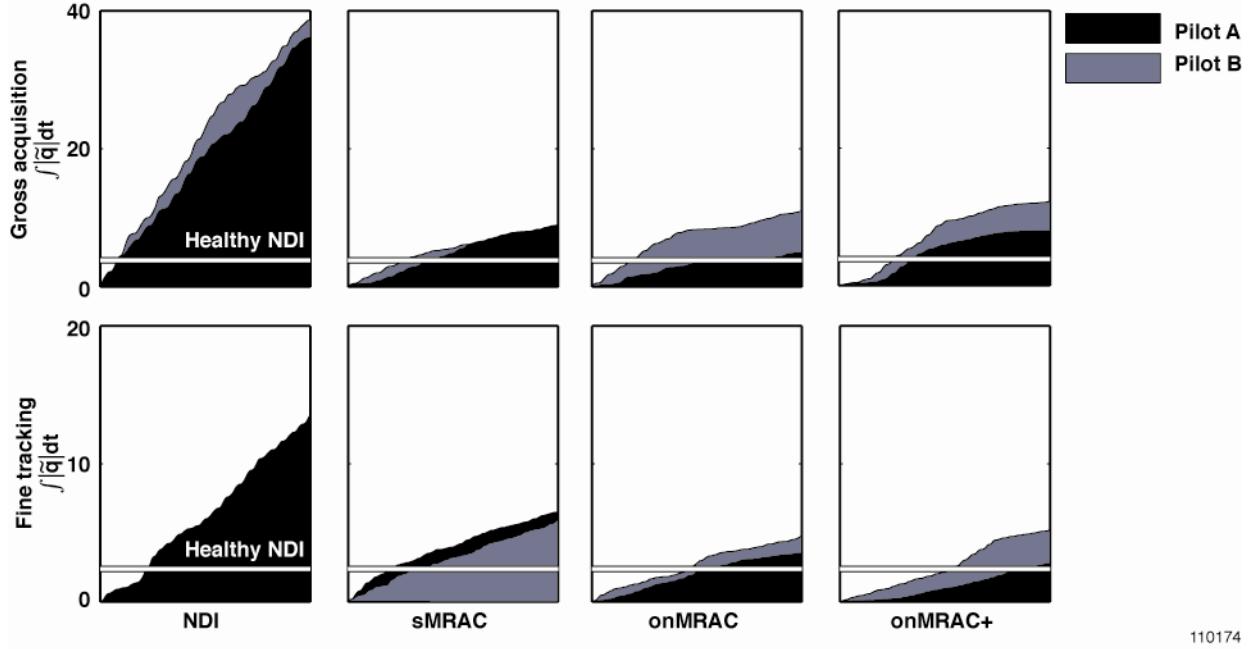


Figure 5. Handling qualities ratings for a 2 g tracking task with a reduced pitch damping failure.

The pitch rate tracking error performance metrics for all four controllers are shown in Fig. 6. The horizontal threshold line indicates the maximum value of this metric for the NDI controller with a healthy aircraft during the same task. This maximum value is essentially the “target score” for the adaptive controllers, indicating whether they have achieved the same level of tracking performance as though the failure had not occurred. For the gross acquisition task, each of the controllers exhibited significantly improved tracking performance when compared to the NDI controller, which is also reflected in the handling qualities rating (HQR) of both pilots from Fig. 5. In the case of fine tracking, the tracking error metric improvements are more modest, which is also reflected in the pilot’s HQR.



110174

Figure 6. Pitch rate tracking error metric for the 2 g tracking task with a reduced pitch damping failure.

Values of the adaptive parameters for the 2 g tracking task with a simulated reduced pitch damping failure are shown in Fig. 7. With no failure in the roll axis, the ideal parameter values for $\hat{\theta}_p$ and $\hat{\sigma}_p$ are close to zero. In the pitch axis, adaptation for each of the controllers is very similar between the two pilots. The roll axis is less similar, primarily due to differences in lateral-directional piloting technique. Pilot A consistently used roll-yaw cross-control during the 2 g tracking task, while Pilot B tended to fly the task without using rudder pedal inputs. The roll axis parameter values for Pilot A are generally larger and more dynamic than those for Pilot B. Both pilots commented on the general lack of predictability for all three adaptive controllers evaluated with this failure.

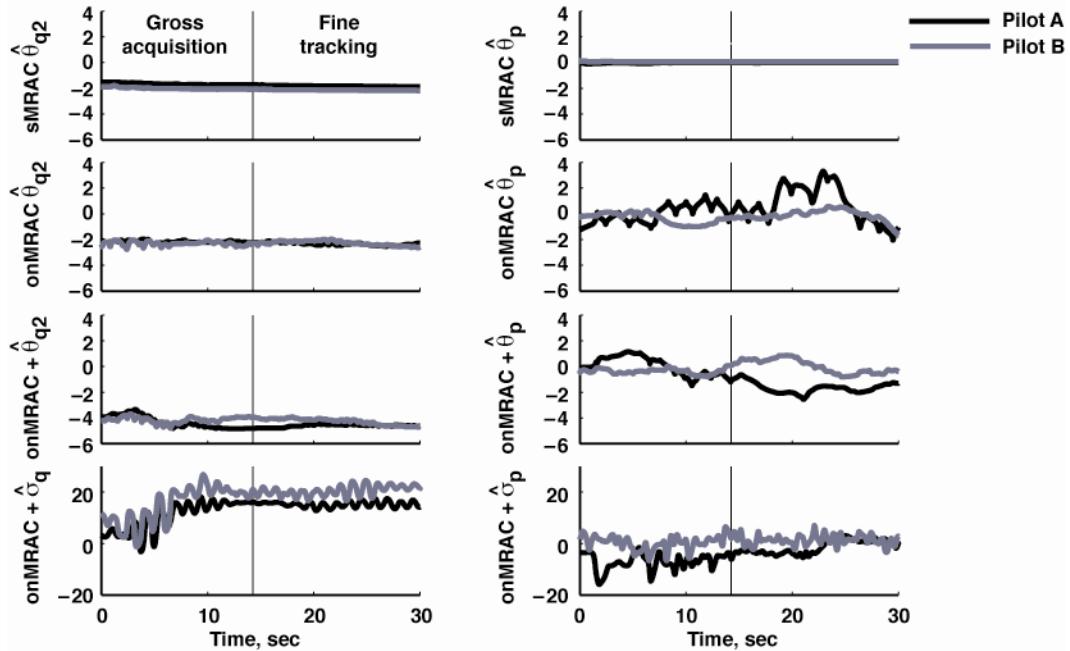


Figure 7. Adaptive parameters during the 2 g tracking task with a reduced pitch damping failure.

1. sMRAC

The adaptive parameters of the sMRAC controller in Fig. 7 can be considered to have converged, in the sense that their values did not change significantly throughout the task. The $\hat{\theta}_p$ parameter is near the expected value of zero. Both pilots commented on the presence of undesirable pitch dynamics, indicating that the effects of the failure were not completely eliminated when $\hat{\theta}_{q2}$ converged to approximately minus two.

2. onMRAC

In the case of the onMRAC controller, $\hat{\theta}_{q2}$ converged to about the same value as that of the sMRAC controller. The tracking error metric for Pilot A in Fig. 6 is consistent with that of the NDI controller with a healthy aircraft indicating that the adaptive controller has largely restored the desired reference model pitch dynamics. However, a comparison of Fig. 2 and Fig. 5 shows that Pilot A's HQR for these scenarios are one entire level apart; the NDI controller with a healthy aircraft received a solid Level 1 rating (an HQR of 2) as shown in Fig. 2, while Fig. 5 shows that the onMRAC with the reduced pitch damping failure was rated as Level 2 (an HQR of 5). Pilot A reported that some undesirable pitch motions were observed during the task accompanied by a “figure-eight motion.”

The presence of noticeable pitch oscillations despite the promising tracking error metric result suggests that this metric alone is not sufficient to adequately evaluate the performance of the adaptive controllers. The observed figure-eight motion is attributable to variations in the roll-axis dynamic response caused by improper adaptation of $\hat{\theta}_p$ illustrated in Fig. 8 for the fine tracking portion of the 2 g tracking task. Because Pilot A flew the 2 g tracking task with roll-yaw cross-control, a semi-steady, non-zero roll rate tracking error developed. Recall from Eq. (28) that $\dot{\hat{\theta}}_p$ is strongly a function of the product of \tilde{x}_p (\tilde{p}) and p . Due to the negative bias in \tilde{p} , $\hat{\theta}_p$ grows with the opposite sign of roll rate. The normalization term of equation Eq. (28) restricts large changes in $\hat{\theta}_p$ to occasions when roll rate is near zero. Consequently, two periods of significant growth in the value of $\hat{\theta}_p$ occur during the fine tracking portion of the task at approximately 18 sec. and again near 22 sec. The ultimate effect is a Level 2 pilot rating due to a destabilization of the roll axis as positive values of $\hat{\theta}_p$ build a positive feedback loop with roll rate. Note that, the OCM term in the update law helps to restrict the growth of $\hat{\theta}_p$ preventing further deterioration of the roll-axis dynamics.

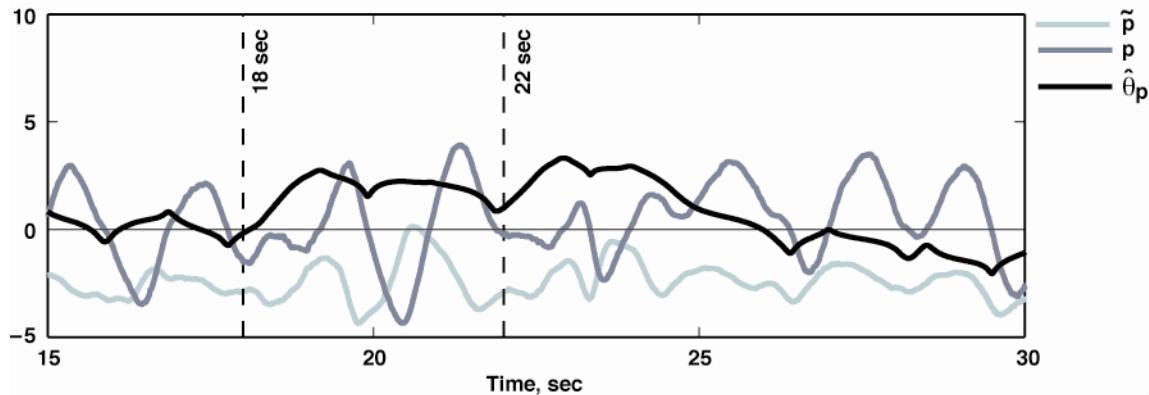


Figure 8. Roll axis onMRAC adaptation during the 2 g fine tracking task with a reduced pitch damping failure (Pilot A).

3. onMRAC+

Pilot A achieved fine tracking performance with the onMRAC+ controller similar to that of the healthy NDI, as measured by the tracking error metric shown in Fig. 6. Pilot B's tracking error metric for the fine tracking portion of the task, as well as for both pilots in the gross acquisition phase, was worse than that for the healthy NDI. Recall from the earlier discussion on the performance of the adaptive controllers with the healthy aircraft that the onMRAC+ controller is prone to adverse pilot-controller interaction, particularly for aggressive inputs. The adaptive parameter $\hat{\sigma}_q$ in Fig. 7 exhibits the tell-tale oscillatory characteristics of this interaction. In general, Pilot B used a more aggressive piloting technique than Pilot A, causing more interaction with the $\hat{\sigma}_q$ parameter and a poorer tracking error metric results, as seen in Fig. 6. Even though he gave all three adaptive controllers identical gross

acquisition CH ratings of 4, Pilot B commented that onMRAC+ was “not my favorite controller” and “I’d love to give this a 4.5, but I have to stick to one of them.”

C. Reduced Roll Damping Failure

Two pilots evaluated the reduced roll damping failure during an in-trail formation tracking task. The ratings of Pilot A and Pilot B for gross acquisition and fine tracking with the simulated roll damping failure are tabulated in Fig. 9. Both pilots rated the NDI controller with no failures (not shown) as Level 1 for this task, with Pilot A giving it CH ratings of 2 and 3, and Pilot B giving it ratings of 3 and 2 for gross acquisition and fine tracking, respectively. Clearly, the failure degraded the gross acquisition handling qualities more than those of the fine tracking portion of the task. All of the adaptive controllers improved the gross acquisition ratings, but the fine tracking results were less conclusive. Generally speaking, the onMRAC+ controller performs as well or better than any of the other controllers, completely restoring the Level 1 fine tracking handling qualities of the NDI controller with no failures.

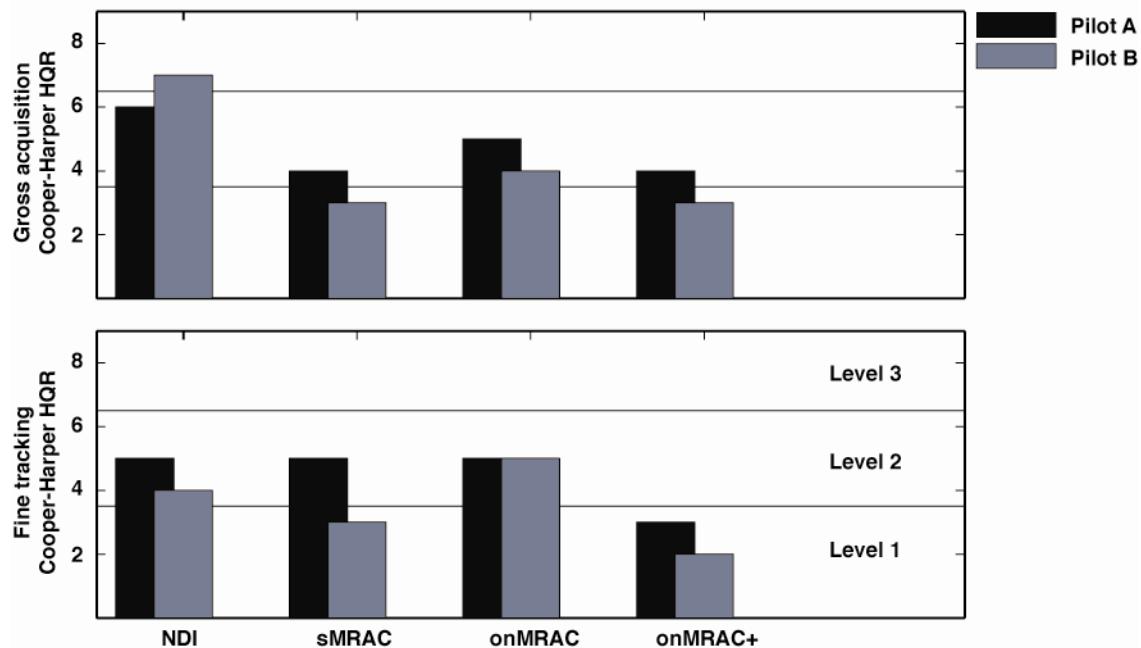


Figure 9. Handling qualities rating for a formation tracking task with a reduced roll damping failure.

Similar to the reduced pitch damping failure, the handling qualities assessments of the roll damping failure have a mixed degree of correlation with the integral of the absolute value of the roll rate tracking error, shown in Fig. 10. All of the adaptive controllers reduce the tracking error as compared to the baseline NDI case. The most significant reduction in this metric occurs for the onMRAC+ fine tracking task, which also received the best ratings.

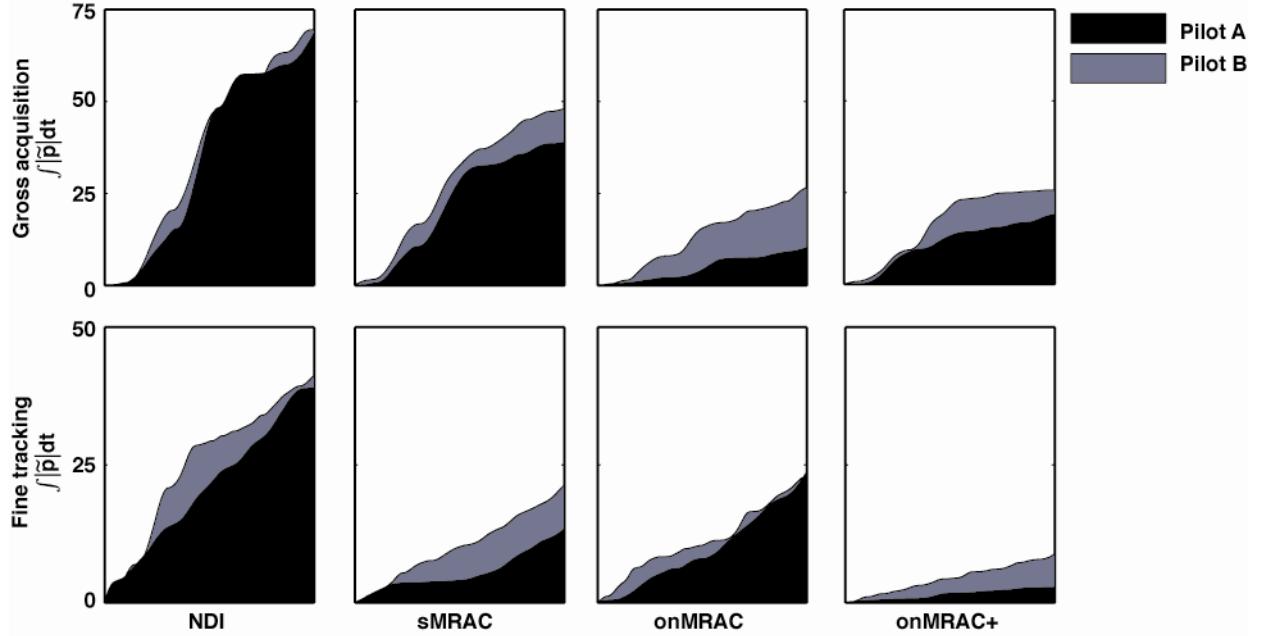


Figure 10. Roll rate tracking error metrics for the formation tracking task with a reduced roll damping failure.

Interestingly, the most improved tracking error metric for the gross acquisition task occurs during Pilot A’s evaluation of the onMRAC controller, which received the worst HQ rating of all the adaptive controllers. An analysis of the data indicates that the lower onMRAC tracking error is attributable to pilot technique. During the sMRAC and onMRAC+ tasks the pilot maneuvered much more aggressively; peak-to-peak roll rate command magnitude ranged between 20% and 50% larger during the gross acquisition task for these two controllers than for the onMRAC controller. It should be noted, that a severe roll axis PIO occurred during Pilot A’s initial fine tracking assessment of the reduced roll damping failure with the onMRAC controller. The task was repeated, and the HQRand tracking error metric shown in Figs. 9 and 10 respectively, are for the second attempt. However, it is possible that the occurrence of the PIO prompted the pilot to perform a less aggressive maneuver during the repeat.

The mechanisms of the roll-axis PIO with the onMRAC controller and a reduced damping failure are very similar to those of the milder roll-axis PIO encountered with the same controller and a reduced pitch damping failure as illustrated in Fig. 8. Both events were triggered by a sustained, quasi-steady roll bias. In the case of the severe PIO during formation tracking with a reduced roll damping failure, the roll bias was due to an asymmetric lift distribution across the wing of the test aircraft caused by the influence of the target aircraft’s wingtip vortex. Due to the geometry of the in-trail tracking task, interference from the wingtip vortex can occur when the nose-to-tail spacing is too small.

The bias in the roll error \tilde{p} can be seen in Fig. 11. From the same figure it is clear that $\hat{\theta}_p$ is initially at the ideal value required to cancel out the failure, but at approximately 142.5 sec. a positive roll rate combines with the positively-biased roll error to drive the adaptive parameter toward zero, effectively re-introducing the failure. Again at about 144.5 sec., the roll rate becomes positive, and $\hat{\theta}_p$ increases to a value equal in magnitude, but opposite in sign to the ideal value essentially doubling the effect of the failure and triggering the PIO.

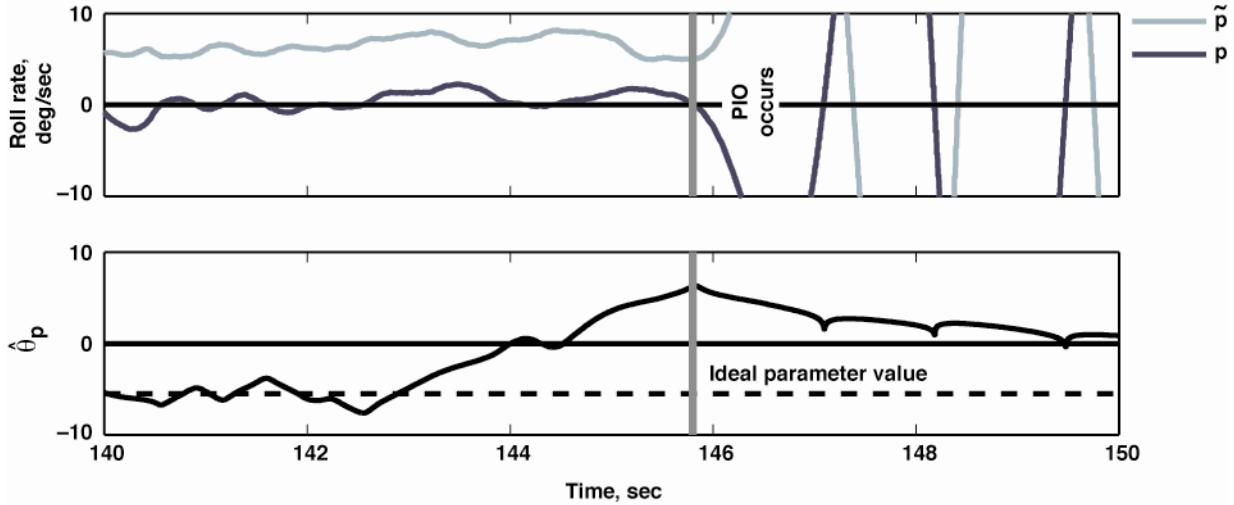


Figure 11. Roll axis onMRAC adaptation during the formation tracking task with a reduced roll damping failure.

VII. Discussion

The research objective was to experimentally determine the performance, handling qualities, and undesirable side-effects of three low-complexity adaptive controllers under a variety of in-flight failure/damage scenarios. This paper deals specifically with matched failures that manifest themselves as changes to the aircraft's damping in either the pitch or roll axis. A companion paper describes the flight results for cross-axis coupling due to simulated damage.¹⁴

Under the assumption stated in this paper's introduction that "simpler is better" with respect to implementation and verification of flight control systems, the following discussion examines the relative strengths and weaknesses of the three adaptive controllers. No recommendation is made as to the appropriate level of complexity, as that is entirely specific to the given application.

A. General Discussion

All three controllers with varying levels of complexity showed improvement in tracking error for reductions in both pitch and roll damping over the baseline non-adaptive NDI controller. Conclusions are more difficult to draw from a comparison between the various adaptive controllers. Using the metric of the integral of the norm of the tracking error, performance was fairly indistinguishable across all three controllers for the reduced pitch damping but appeared to improve with increasing complexity for the reduced roll damping failure.

Handling qualities ratings for the three adaptive controllers under simulated failure scenarios did not differ significantly, although the onMRAC+ controller consistently had ratings as good as or better than the other two. This trend does not extend to the healthy aircraft case, in which the onMRAC+ has appreciably degraded handling qualities.

B. sMRAC Discussion

The sMRAC controller provided a small measure of performance and handling qualities improvement for both simulated failure scenarios and displayed no significant negative side effects. For the healthy aircraft case, this adaptive controller degraded the error tracking performance slightly but retained Level 1 handling qualities. The sMRAC adaptive parameters shown in Fig. 4 are typical of tests conducted with this controller showing very little variation in parameter value with dynamic maneuvering.

C. onMRAC Discussion

The onMRAC controller performed well in the healthy aircraft configuration but displayed a tendency to be PIO prone. The undesirable roll axis adaptation in the presence of biased roll commands can be corrected by applying a high-pass filter to all commands and feedbacks to the roll adaptive controller. This approach was used successfully in the pitch axis to prevent similar poor adaptation characteristics during extended pitch maneuvers, such as wind-up turns.

D. onMRAC+ Discussion

A drawback of the onMRAC+ controller is its proclivity for adverse interactions with the pilot during aggressive pitch inputs. The activation of the OCM term on the $\hat{\sigma}_q$ adaptive parameter has shown promise through simulation in preventing these adverse interactions. Correcting this negative side effect of the controller should improve its HQ in the healthy aircraft case, and possibly in the various failure scenarios as well. The improvement in performance achieved with this controller may justify the slight increase in complexity over the other controllers that were evaluated depending upon the specific application.

VIII. Summary

Flight research was conducted at the NASA Dryden Flight Research Center to study the effects of varying levels of adaptive controller complexity on the performance and handling qualities of an aircraft under various simulated failure or damage conditions. Four controllers were designed with varying levels of complexity including a baseline, non-adaptive nonlinear dynamic inversion control law. The remaining three were model reference adaptive control designs, the simplest of which consisted of a single adaptive parameter in each of the pitch and roll axes computed using a basic gradient-based update law. A second design was built upon the first by increasing the complexity of the update law. The third and most complex design added an additional adaptive parameter to each axis.

The controllers showed improvement in reference model tracking for simulated failures involving severe changes to the pitch and roll damping of the airplane. Handling qualities evaluations of the controllers revealed that despite their improved tracking, the adaptive controllers occasionally exhibited undesirable side-effects of unpredictability and increased pilot workload. Recommendations were given for modification to the controllers that may eliminate these side effects and bring their performance and handling qualities in line with that of the healthy aircraft, even in the presence of failures.

References

- ¹Whitaker, H. P., "Massachusetts Institute of Technology Presentation," *Proceedings of the Self Adaptive Flight Control Systems Symposium*, Wright Air Development Center, 13-14 January 1959, pp. 58-78.
- ²Schuck, H. O., "Honeywell's History and Philosophy in the Adaptive Control Field," *Proceedings of the Self Adaptive Flight Control Systems Symposium*, Wright Air Development Center, 13-14 January 1959, pp. 123-145.
- ³NASA Flight Research Center Staff, "Experience with the X-15 Adaptive Flight Control System," NASA TN D-6208, 1971.
- ⁴Thompson, M. O., and Welsh, J. R., "Flight Test Experience with Adaptive Control Systems," Advanced Control System Concepts, AGARD CP No. 58, *AGARD Joint Symposium of the Guidance and Control Panel, and the Flight Mechanics Panel of AGARD*, Oslo Norway, 3-5 September 1968, pp. 141-147, 1970.
- ⁵Monaco, J., Ward, D., Barron, R., and Bird, R., "Implementation and Flight Test Assessment of an Adaptive, Reconfigurable Flight Control System," AIAA-97-3738, 1997.
- ⁶Bosworth, J. T., "Flight Results of the NF-15B Intelligent Flight Control System (IFCS) Aircraft with Adaptation to a Longitudinally Destabilized Plant," AIAA-2008-6985.
- ⁷Burken, J. J., Hanson, C. E., Lee, J. A., and Kaneshige, J. T., "Flight Test Comparison of Different Adaptive Augmentations to Fault Tolerant Control Laws for a Modified F-15 Aircraft," AIAA-2009-2056.
- ⁸Krishnakumar, K., "A Request for Information in Flight Validation of Adaptive Control to Prevent Loss-of-Control Events," Solicitation Number: NNH09ZEA002L, NASA Headquarters Office of Procurement, April 2009.
- ⁹Ioannou, P.A., and Sun, J., *Robust Adaptive Control*, Prentice Hall, Inc., Upper Saddle River, New Jersey, 1996.
- ¹⁰Greenwood, D.T., *Principles of Dynamics*, Prentice Hall, Inc., Upper Saddle River, New Jersey, 1988.
- ¹¹Miller, C. J., "Nonlinear Dynamic Inversion Baseline Control Law: Architecture and Performance Predictions," AIAA-2011-#### (to be published), 2011.
- ¹²McRuer, D., Ashkenas, I., and Graham, D., *Aircraft Dynamics and Automatic Control*, Princeton University Press, Princeton, New Jersey, 1973.
- ¹³Nguyen, N., "Asymptotic Linearity of Optimal Control Modification Adaptive Law with Analytical Stability Margins," AIAA-2010-3301.
- ¹⁴Schaefer, J., Hanson, C., Johnson, M., and Nguyen, N., "Handling Qualities for Model Reference Adaptive Controllers of Varying Complexity with Pitch-Roll Coupled Failures," AIAA-2011-#### (to be published), 2011.
- ¹⁵Gray, W., "Handling Qualities Evaluation at the USAF Test Pilot School," AIAA-2009-6317, 2009.
- ¹⁶Cooper, G. E., and Harper, R. P., "The Use of Pilot Rating in the Evaluation of Aircraft Handling Qualities," NASA TN D-5153, April 1969.
- ¹⁷Stepanyan, V., Krishnakumar, K., Nguyen, N., and Van Eykeren, L., "Stability and Performance Metrics for Adaptive Flight Control," AIAA 2009-5965, 2009.

¹⁸Miller, C. J., "Nonlinear Dynamic Inversion Baseline Control Law: Flight Test Results for FAST F/A 18," AIAA-2011-#### (to be published), 2011.

¹⁹Hanson, C., Johnson, M., and Schaefer, J., "Design of a Low Complexity Model Reference Adaptive Controller," NASA/TM-####-215972 (to be published).